Towards the Integration of an Intuitionistic First-Order Prover into Coq

Fabian Kunze
Saarland University
s9fakunz@stud.uni-saarland.de

An efficient intuitionistic first-order prover integrated into Coq is useful to replay proofs found by external automated theorem provers. We propose a two-phase approach: An intuitionistic prover generates a certificate based on the matrix characterization of intuitionistic first-order logic; the certificate is then translated into a sequent-style proof.

1 Introduction

Sledgehammer [11] and HOLyHammer [5] drastically improved the productivity for users of proof assistants. They make the capabilities of automated theorem provers (ATPs) available from within interactive proof assistants.

The large, monolithic design of state-of-the-art theorem provers can not be easily trusted to be free of bugs. Thus invoking theorem provers as an oracle is unacceptable for most users. Proof assistants are more trustworthy because all reasoning is checked by a kernel intentionally kept small.

To integrate external provers, small yet efficient, *certified* provers *integrated* into the proof assistant are used: Although it is often possible to mechanically translate the proof to a format accepted by the proof assistant, the integrated prover allows for the reconstruction without the full knowledge of all axioms and rules used by the external prover. Thus an integrated prover simplifies the integration of not only one but different external provers.

There has been effort to integrate classical provers into Coq, e.g. SMTCoq [1], Satallax [3] and why3 [2], but they produce proofs that assume classical axioms. As a fair amount of proof developments avoids assuming additional axioms, the acceptance of a future 'Coq Hammer' benefits from the integration of an efficient, *intuitionistic* prover.

2 Existing Intuitionistic Provers in Coq

The existing intuitionistic first-order provers integrated into Coq are not very strong. We evaluated `firstorder` [4], a built-in tactic based on a sequent calculus, and JProver [14], a plugin available for Coq. Using Coq version `8.6pl1`, we considered first-order problems that are likely to emerge in a future 'Coq Hammer'.

For example, we tested formulas where the instantiate of quantifiers is not immediately determined using a goal-driven approach:

$$\begin{aligned}&(\forall x, x = x) \wedge (\forall x, Px \vee Qx)\\ \wedge\ &(\forall xy, x = y \wedge Px \Rightarrow Ry) \wedge (\forall xy, x = y \wedge Qx \Rightarrow Ry) \Rightarrow (\forall x, Rx).\end{aligned}$$

J.C. Blanchette and C. Kaliszyk:
Hammers for Type Theories (HaTT'16)
EPTCS 210, 2016, pp. 30–35, doi:10.4204/EPTCS.210.6

© Fabian Kunze

On this formula, `firstorder` was unable to find a proof during the several minutes we run it. JProver succeeded in less than one second.

We also invoked both provers on several set-theoretical problems from the ILTP (Intuitionistic Logic Theorem Proving) library [12]. Similar to the intended use case, we only supplied the axioms needed for the proofs, resulting in problems like

$$\begin{aligned}&(\forall ABX, X \in A \cup B \Leftrightarrow X \in A \vee X \in B)\\ \wedge\ &(\forall AB, A = B \Leftrightarrow A \subseteq B \wedge B \subseteq A)\\ \wedge\ &(\forall AB, A \subseteq B \Leftrightarrow \forall X, X \in A \Rightarrow X \in B) \Rightarrow (\forall A, A \cup A = A).\end{aligned}$$

On this and similar problems, both `firstorder` and JProver failed to find proofs before we aborted them after running several minutes.

Therefore, faster intuitionistic provers integrated into Coq are necessary for a 'Coq Hammer' used in practice.

3 Proposed Architecture

We propose to employ the recent improvements on automated, intuitionistic first-order theorem proving by Otten: ileanCoP [7, 8] and the forthcoming intuitionistic version of nanoCoP [10, 9]. The existing implementations of both provers verified that the formulas in Section 2 are valid in under a second. Both provers are based on the existence of proof certificates for the matrix characterization of (intuitionistic) validity [15], which can be translated to sequent-style proofs [13].

This architecture is similar to that of JProver (which uses the same characterization of validity), but uses a more efficient proof search procedure, leading to a higher success rate.

3.1 Finding Proof Certificates

The performance of ileanCoP is well in identifying valid formulas, compared to other intuitionistic provers [8]. But it does not keep track of the proof found. Furthermore, it is based on a *clausal* variant of the matrix characterization for intuitionistic logic. The necessary translation into a non-clausal matrix proof has been sketched in the correctness proof of ileanCoP [7], but to our knowledge has not yet been implemented.

The classical prover nanoCoP [10] solves both problems: It outputs the proof certificate found and uses the non-clausal matrix characterization of classical validity. Otten is currently extending nanoCoP to an intuitionistic variant by integrating prefix unification [15], a method already employed to derive ileanCoP from the classical prover leanCoP.

In our proposed architecture, the proof certificate for a first-order formula F consists of a *multiplicity* μ, a pair of substitutions $\sigma = (\sigma_Q, \sigma_J)$ and a set of pairs of σ*-complementary* atoms in the formula (connections) that *spans* F^μ.

We will now give an very informal intuition about this certificate.

A Part of the certificate is already needed for the matrix characterization of classical logic: The multiplicity μ takes care of the multiple instances an all-quantified subformula of F may be needed in the proof. One part of 'σ-complementary' ensures that that two atoms in a connection are identical under the (non-circular) term substitution σ_Q, but have different polarity.

The set of connections *spans* the formula if every *path* through the formula contains at least one connections. In the quantifier-free case, each path correspond to a disjunction in the conjunctive normal

form. In the case of formulas with quantifiers, each path correspond a branch of a (classical) analytic tableaux, where quantifiers are instantiated according to σ_Q.

The main difference in the intuitionistic characterization is the use of σ_J to ensure that the positions of the pair of complementary atoms in the formula are 'compatible'. The position of an atom is defined by structural recursion on the formula and represented by a string, consisting of fresh constants and fresh variables.

An example of this for an intuitionistic valid formula is $P \Rightarrow P$, where the two atoms can be made complementary: The position of the first P is described by the string z with a fresh variable z, while the position of the second P is the string a consisting of a fresh constant a. Defining $\sigma_J(z) = a$ unifies those strings.

For the formula $\neg P \vee P$, a theorem of classical, but not intuitionistic logic, the two atoms can not be made complementary: The position of the first P is described by xa, while the one for the second P is b. As the second position contains no variable and no a, we can never unify those strings.

This concept generalizes to quantified formulas, but for the main idea, it suffices to study the cases for non-quantified formulas.

For a more formal definition and a few more examples, we recommend the first two Sections of [7], and Chapter 8, §4 of [15].

It should be noted that one of the main improvements of nanoCoP compared to JProver is the handling of the multiplicities: nanoCoP adds instances of subformulas during the proof search as needed, while JProver fixes the multiplicity before searching for an proof; on failure, an additional instance of the whole formula gets added and the proof is retried. Although both are complete, the first approach is more goal-driven and thus expected to be more efficient.

3.2 Generating Sequence Proofs

The high-level idea is that the proof certificate guarantees that on each branch of the sequent-style proof, eventually complementary atoms are found. The difficulty is to traverse the formulas in the right order, which depends on σ_J.

The translation of a matrix characterisation proof certificate into a sequent-style proof has already been investigated and implemented for JProver[13]. We intend to adopt this translation, as we expect it to be reasonable fast: In the examples we tried and where JProver succeeded, the sequence-style proof produced was rather short. In the cases where JProver did not succeed in an acceptable time, it did not even reach the sequence-proof generation. Thus we conclude that the bottleneck of JProver, at least in the examples we tried, is the proof certificate search.

4 Discussion

Modular vs Monolithic

We explicitly want to use a modular implementation for the two phases, possibly written in multiple languages. The Prolog version of the intuitionistic variant of nanoCoP is expected to materialize soon and there is already an implementation of the sequence proof generating algorithm integrated into Coq. Thus we expect no challenge in creating a prototype of the suggested architecture using the Prolog program. This would allow us to test whether the proposed setup is suitable for the intended use case.

In the longer term, it would be desirable to have a native OCaml implementation of the proof search procedure, allowing for a deployment within Coq, without additional binaries. The classical leanCoP

has been ported to OCaml for the HOL light proof assistant, with performence comparable to the Prolog version[6]. This port can serve as a starting point for a native OCaml version of the forthcoming intuitionistic nanoCoP. Then, the modular approach allows to optionally use external proof procedures. This allows to evaluate improvements to the Prolog proof procedure before porting them.

Also, a modular design allows to more easily use parts an implementation this for other, intuitionistic proof assistants. This additional usage should be kept in mind while developing this, and other, tools towards Hammers in Type Theory.

Explicit Proofs vs Reflection

One approach in proof automation in Coq is 'proof by reflection': Some or all parts of the the proof search procedure are written in Coq, including a correctness proof. The proof of a statement then *is* the call to this Coq procedure.

One argument for 'proof by reflection' in Coq is the efficiency. But this is just a benefit compared to an implementation using Ltac, the tactic language in Coq: The evaluation of native Coq terms is heavily optimized to the extend of native machine code compilation and execution. In contrast, Ltac is just interpreted on top of several layers of abstraction. As we propose to use OCaml, not Ltac, for the computationally intense parts, this argument does not apply here.

We assume that the search for the proof certificate could be more easily written, modified, or enriched with heuristics, when using a language allowing side effect. This discourages the use of reflection in the first part of our proposed architecture.

Reflection seems to be more reasonable for the second part, the translation to a sequence proof: There is no need to explicitly generate the sequence proof when a certified procedure guarantees that the sequence proof *does* exist when the certificate satisfies the appropriate conditions.

The challenge here would be that the proof certificate must annotated with type information rich enough to reduce to proofs for all formulas we intend to proof: This means that when the terms in the formula are not single sorted, but have of more complex types, e.g. dependent types, this must be incorporated in the proof certificate, the translation procedure itself and its correctness proof. At first, it seems that a benefit would be that the translation is proven to be sound by design. But to check the conditions that a proof certificate is indeed valid is more or less computationally equivalent hard as to generating a sequence-style proof.

Another aspect to consider is that some usage, a formula that is not first-order can be transformed into an first-order formula such that a proof of the later formula can be translated back to a proof of the former formula. In a reflective proof reconstruction, this intermediate steps may can not type-check.

Intuitionistic vs Classical

Automated theorem proving in intuitionistic logic is computationally harder than in classical logic. For developments assuming classical axioms, the intuitionistic part of both phases can be made optional, resembling the classical proof search of nanoCoP without significant overhead.

Note that the proof search in this proposed architecture does neither need skolemization nor clausal normal forms. Thus more structure of the different lemmas and parts of the formulas is preserved and in some sense, this approach is closer to humans reasoning. Further investigation of this architecture could lead to insights useful for automated reasoning in proof assistants of classical logic.

Acknowledgements

We thank Jens Otten for his helpful discussions and suggestions, and Jasmin Blanchette and the anonymous reviewers for their comments on this extended abstract.

References

[1] Michael Armand, Germain Faure, Benjamin Grégoire, Chantal Keller, Laurent Théry & Benjamin Werner (2011): *Certified Programs and Proofs: First International Conference, CPP 2011, Kenting, Taiwan, December 7-9, 2011. Proceedings*, chapter A Modular Integration of SAT/SMT Solvers to Coq through Proof Witnesses, pp. 135–150. Springer Berlin Heidelberg, Berlin, Heidelberg. Available at http://dx.doi.org/10.1007/978-3-642-25379-9_12.

[2] François Bobot, Jean-Christophe Filliâtre, Claude Marché & Andrei Paskevich (2011): *Why3: Shepherd Your Herd of Provers*. In K. Rustan M. Leino & Michał Moskal, editors: *Boogie 2011*, pp. 53–64.

[3] Chad E. Brown (2012): *Satallax: An Automatic Higher-Order Prover*. In Bernhard Gramlich, Dale Miller & Uli Sattler, editors: *Automated Reasoning—6th International Joint Conference, IJCAR 2012, Manchester, UK, June 26-29, 2012. Proceedings*, *Lecture Notes in Computer Science* 7364, Springer, pp. 111–117. Available at http://dx.doi.org/10.1007/978-3-642-31365-3_11.

[4] Pierre Corbineau (2004): *First-Order Reasoning in the Calculus of Inductive Constructions*, pp. 162–177. Springer Berlin Heidelberg, Berlin, Heidelberg. Available at http://dx.doi.org/10.1007/978-3-540-24849-1_11.

[5] Cezary Kaliszyk & Josef Urban (2015): *HOL(y)Hammer: Online ATP Service for HOL Light*. *Mathematics in Computer Science* 9(1), pp. 5–22. Available at http://dx.doi.org/10.1007/s11786-014-0182-0.

[6] Cezary Kaliszyk, Josef Urban & Jiři Vyskočil (2015): *Certified Connection Tableaux Proofs for HOL Light and TPTP*. In: *Proceedings of the 2015 Conference on Certified Programs and Proofs*, CPP '15, ACM, New York, NY, USA, pp. 59–66. Available at http://doi.acm.org/10.1145/2676724.2693176.

[7] Jens Otten (2005): *Clausal Connection-Based Theorem Proving in Intuitionistic First-Order Logic*. In Bernhard Beckert, editor: *Automated Reasoning with Analytic Tableaux and Related Methods: 14th International Conference, TABLEAUX 2005, Koblenz, Germany, September 14-17, 2005. Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 245–261. Available at http://dx.doi.org/10.1007/11554554_19.

[8] Jens Otten (2008): *leanCoP 2.0 and ileanCoP 1.2: High Performance Lean Theorem Proving in Classical and Intuitionistic Logic (System Descriptions)*. In Alessandro Armando, Peter Baumgartner & Gilles Dowek, editors: *Automated Reasoning: 4th International Joint Conference, IJCAR 2008 Sydney, Australia, August 12-15, 2008 Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 283–291. Available at http://dx.doi.org/10.1007/978-3-540-71070-7_23.

[9] Jens Otten (2016): personal communication.

[10] Jens Otten (2016): *nanoCoP: A Non-clausal Connection Prover*. In: *Automated Reasoning: 8th International Joint Conference, IJCAR 2016 Coimbra, Portugal, June 27 -July 2, 2016 Proceedings*. To appear.

[11] Lawrence C. Paulson & Jasmin Christian Blanchette (2010): *Three years of experience with Sledgehammer, a Practical Link Between Automatic and Interactive Theorem Provers*. In Geoff Sutcliffe, Stephan Schulz & Eugenia Ternovska, editors: *The 8th International Workshop on the Implementation of Logics, IWIL 2010, Yogyakarta, Indonesia, October 9, 2011*, *EPiC Series* 2, EasyChair, pp. 1–11. Available at http://www.easychair.org/publications/?page=820355915.

[12] Thomas Raths, Jens Otten & Christoph Kreitz (2007): *The ILTP Problem Library for Intuitionistic Logic: Release v1.1*. *Journal of Automated Reasoning* 38(1-3), pp. 261–271. Available at http://dx.doi.org/10.1007/s10817-006-9060-z.

[13] Stephan Schmitt & Christoph Kreitz (1996): *Converting non-classical matrix proofs into sequent-style systems*. In: *Automated Deduction — Cade-13: 13th International Conference on Automated Deduction New Brunswick, NJ, USA, July 30 – August 3, 1996 Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 418–432. Available at `http://dx.doi.org/10.1007/3-540-61511-3_104`.

[14] Stephan Schmitt, Lori Lorigo, Christoph Kreitz & Alexey Nogin (2001): JProver *: Integrating Connection-based Theorem Proving into Interactive Proof Assistants*. In R. Gore, A. Leitsch & T. Nipkow, editors: *International Joint Conference on Automated Reasoning, Lecture Notes in Artificial Intelligence* 2083, Springer Verlag, pp. 421–426, doi:10.1007/3-540-45744-5_34.

[15] Lincoln Wallen (1990): *Automated Deduction in Nonclassical Logics*. MIT Press, Cambridge, MA, USA.